

## AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

### Listing of Claims:

1-26. (Cancelled)

27. (Previously presented) A method of parallelizing an operation, the method comprising the steps of:

dividing the operation into a set of work partitions;

assigning work partitions from said set of work partitions to a plurality of entities,

wherein at least one entity of said plurality of entities is assigned a plurality of

work partitions from said set of work partitions; wherein the step of assigning

work partitions is performed by assigning the work partitions in a sequence based

at least in part on sizes associated with the work partitions; and

said plurality of entities operating in parallel on work partitions assigned to them to

perform said operation.

28. (Previously presented) The method of Claim 27 wherein the step of assigning the work partitions in a sequence is performed by assigning relatively larger work partitions before assigning relatively smaller work partitions.

1 29. (Previously presented) A method of parallelizing an operation, the method comprising the  
2 steps of:  
3 dividing the operation into a set of work partitions;  
4 assigning work partitions from said set of work partitions to a plurality of entities,  
5 wherein at least one entity of said plurality of entities is assigned a plurality of  
6 work partitions from said set of work partitions, wherein the step of assigning  
7 work partitions includes:  
8 assigning said at least one entity a first work partition from said set of work  
9 partitions; and  
10 after said at least one entity has completed operation on said first work partition,  
11 assigning said at least one entity a second work partition from said set of work  
12 partitions, wherein the step of assigning said at least one entity a second work  
13 partition includes  
14 determining whether there are any unassigned work partitions from a first level in  
15 the hierarchy to which said first work partition belonged; and  
16 if there are no unassigned work partitions from the first level in the  
17 hierarchy, then selecting said second work partition from a level in  
18 said hierarchy that is two levels above said first level in said  
19 hierarchy;  
20 said plurality of entities operating in parallel on work partitions assigned to them to  
21 perform said operation; and  
22 wherein the operation is specified in a query that corresponds to a hierarchy of  
23 operations.

1 30. (Previously presented) A method of parallelizing an operation, the method comprising the  
2 steps of:  
3 dividing the operation into a set of work partitions;  
4 assigning work partitions from said set of work partitions to a plurality of entities,  
5 wherein at least one entity of said plurality of entities is assigned a plurality of  
6 work partitions from said set of work partitions;  
7 said plurality of entities operating in parallel on work partitions assigned to them to  
8 perform said operation;  
9 the method includes the step of generating a serial execution plan for operations in a  
10 database management system (DBMS) running on a computer system;  
11 the method includes the step of generating a parallelized execution plan for said serial  
12 execution plan, said parallelized execution plan including first and second  
13 operations;  
14 the step of dividing an operation is performed by dividing said second operation;  
15 the plurality of entities includes one or more slave processes operating on a plurality of  
16 data partitions, the quantity of said data partitions being greater than the quantity  
17 of said slave processes;  
18 executing said parallelized execution plan when a plurality of parallel resources of said  
19 computer system are available; and  
20 executing said serial execution plan when said plurality of resources are not available.

1 31. (Previously presented) The method of claim 30 wherein said step of generating a  
2 parallelized execution plan includes the steps of:

3 identifying one or more segments of said serial execution plan that can be parallelized;

4 and

5 identifying partitioning requirements of said one or more segments.

1 32. (Previously presented) The method of claim 30 wherein said step of generating a  
2 parallelized execution plan is based on a specification of parallelism in a statement  
3 specifying one of said operations.

1 33. (Previously presented) A method of parallelizing an operation, the method comprising the  
2 steps of:

3 dividing the operation into a set of work partitions;

4 assigning work partitions from said set of work partitions to a plurality of entities,

5 wherein at least one entity of said plurality of entities is assigned a plurality of

6 work partitions from said set of work partitions;

7 said plurality of entities operating in parallel on work partitions assigned to them to

8 perform said operation;

9 generating an execution plan for said operation;

10 examining said execution plan from bottom up;

11 identifying a parallelized portion of said execution plan, said parallelized portion can be

12 processed in parallel, said parallelized portion including first and second

13 operations, said first and second operations being executable in parallel;

14 wherein the step of dividing the operation is performed by dividing said second

15 operation;

16 wherein the plurality of entities includes one or more slave processes operating on a  
17 plurality of data partitions, the quantity of said data partitions being greater than  
18 the quantity of said slave processes;  
19 identifying some serial portion of said execution plan, said serial portion can be  
20 processed in serial; and  
21 allocating a central scheduler between said parallelized portion and said serial portion.

1 34. (Previously presented) The method of Claim 33 further including the steps of:  
2 identifying a first data flow requirement for a first portion of said execution plan said first  
3 data flow requirement corresponding to a partitioning of a data flow required by  
4 said first portion;  
5 identifying a second data flow requirement for a second portion of said execution plan  
6 said second data flow requirement corresponding by said second portion; and  
7 allocating a data flow director between said first portion and said second portion when  
8 said first data flow requirement is not compatible with said second data flow  
9 requirement said data flow director repartitioning a data flow of said first portion  
10 to be compatible with said second data flow requirement.

1 35. (Previously presented) A method for parallelizing an operation, the method comprising  
2 the steps of:  
3 dividing the operation into a set of work partitions;

4 assigning work partitions from said set of work partitions to a plurality of entities,  
5 wherein at least one entity of said plurality of entities is assigned a plurality of  
6 work partitions from said set of work partitions;  
7 said plurality of entities operating in parallel on work partitions assigned to them to  
8 perform said operation;  
9 generating an execution plan to execute database management system (DBMS)  
10 operations in parallel, said execution plan including first and second operations;  
11 wherein the step of dividing said operation is performed by dividing said second  
12 operation;  
13 initiating an operation coordinator in a computer system to coordinate execution of said  
14 execution plan;  
15 initiating, by said operation coordinator, a first set of slaves operating on a plurality of  
16 data partitions to produce data, the quantity of said data partitions being greater  
17 than the quantity of said first set of slave processes;  
18 initiating, as said plurality of entities, by said operation coordinator, a second set of  
19 slaves to consume data; and  
20 directing said second set of slaves to produce data and said first set of slaves to consume  
21 data when said first set of slaves finishes producing data.

1 36. (Previously presented) The method of claim 35 wherein said execution plan is comprised  
2 of operator nodes and said operator nodes are linked together to form execution sets.

1 37. (Previously presented) A method for parallelizing an operation, the method comprising  
2 the steps of:  
3 dividing the operation into a set of work partitions;  
4 assigning work partitions from said set of work partitions to a plurality of entities,  
5 wherein at least one entity of said plurality of entities is assigned a plurality of  
6 work partitions from said set of work partitions;  
7 said plurality of entities operating in parallel on work partitions assigned to them to  
8 perform said operation;  
9 generating an execution plan to execute said operations in parallel, said execution plan  
10 including first and second operations;  
11 wherein the step of dividing said operation includes dividing said first operation;  
12 initiating producer slaves operating on a plurality of data partitions to produce a first data  
13 production;  
14 initiating consumer slaves to consume said first data production;  
15 when said first data production is completed, generating an identification of a plurality of  
16 said consumer slaves that did not receive data in said first data production;  
17 examining said identification during a subsequent data production; and  
18 reducing said subsequent data production such that said subsequent data production does  
19 not produce data for said plurality of said consumer slaves.

1 38. (Currently Amended) A method for processing a query, the method comprising the steps  
2 of:  
3 receiving a statement that specifies at least an operation;

4 determining a user-specified degree of parallelism to use in performing the operation;  
5 dividing the operation into a set of work partitions;  
6 performing a determination of how many entities to use to perform said operation based,  
7 at least in part, on the user-specified degree of parallelism;  
8 assigning work partitions from said set of work partitions to a plurality of entities based  
9 on said determination; and  
10 said plurality of entities operating in parallel on work partitions assigned to them to  
11 perform said operation.

1 39. (Currently amended) The method of Claim 38 wherein:  
2 the query requires a plurality of operations;  
3 the user-specified degree of parallelism is specified in said statement, and  
4 the statement specifies said degree of parallelism for a subset of the plurality of  
5 operations required by the query.

1 40. (Currently amended) The method of Claim 39 wherein  
2 the user-specified degree of parallelism is specified in said statement; and  
3 the degree of parallelism specified by the query indicates that no amount of parallelism is  
4 to be used during execution of a particular portion of the query.

1 41. (Currently amended) The method of Claim 38 wherein  
2 the user-specified degree of parallelism is specified in said statement, and



3 the degree of parallelism specified by the query indicates a maximum amount of  
4 parallelism to use during execution of said operation.

1 42. (Previously presented) A method of processing a query, the method comprising the steps  
2 of:

3 dividing an operation required by said query into a set of work partitions by generating a  
4 set of query fragments;

5 incorporating hints into at least some of said query fragments, wherein the hint associated  
6 with a given query fragment indicates how to perform the work partition

7 associated with said given query fragment;

8 assigning query fragments from said set of query fragments to a plurality of entities; and

9 said plurality of entities operating in parallel on query fragments assigned to them to

10 perform said operation, wherein entities working on a query fragment associated

11 with a hint perform the work partition associated with said query fragment in a

12 manner dictated by said hint.

1 43. (Previously presented) The method of Claim 42 wherein the step of incorporating hints

2 includes incorporating hints that dictate the operation of a table scan.

1 44. (Previously presented) The method of Claim 43 wherein the step of incorporating hints

2 that dictate the operation of a table scan includes incorporating hints that rowid

3 partitioning is to be used during the table scan.

- 1 45. (Previously presented) The method of Claim 42 wherein the step of incorporating hints  
2 includes incorporating hints that specify performance of a full table scan.
- 1 46. (Previously presented) The method of Claim 42 wherein the step of incorporating hints  
2 includes incorporating hints that specify using a particular type of join.
- 1 47. (Previously presented) The method of Claim 46 wherein the step of incorporating hints  
2 that specify using a particular type of join includes incorporating hints that specify using  
3 a sort/merge join.
- 1 48. (Previously presented) The method of Claim 46 wherein the step of incorporating hints  
2 that specify using a particular type of join includes incorporating hints that specify using  
3 a nested loop join.
- 1 49. (Previously presented) A method of processing a query, the method comprising the steps  
2 of:  
3 determining a hierarchy of operations associated with a query;  
4 dividing a first operation required by said query into a first set of work partitions;  
5 dividing a second operation required by said query into a second set of work partitions,  
6 wherein said second operation immediately follows said first operation in said  
7 hierarchy;

dividing a third operation required by said query into a third set of work partitions,  
wherein said third operation immediately follows said second operation in said  
hierarchy;  
assigning work partitions from said first set of work partitions to a first plurality of  
entities;  
said first plurality of entities operating in parallel on work partitions assigned to them  
from said first set of work partitions to perform said first operation;  
assigning work partitions from said second set of work partitions to a second plurality of  
entities, wherein said second plurality of entities are different entities than said  
first plurality of entities; and  
said second plurality of entities operating in parallel on work partitions assigned to them  
from said second set of work partitions to perform said second operation;  
assigning work partitions from said third set of work partitions to said first plurality of  
entities; and  
said first plurality of entities operating in parallel on work partitions assigned to them  
from said third set of work partitions to perform said third operation.

50. (Previously presented) The method of Claim 49 further comprising performing the  
following steps when a given entity in said first set of entities finishes performing a work  
partition from said first set of work partitions:  
determining whether there are any unassigned work partitions from said first set of work  
partitions; and

6 if there are no unassigned work partitions from said first set of work partitions, then  
7 assigning the given entity a work partition selected from said third set of work  
8 partitions; and  
9 if there are unassigned work partitions from said first set of work partitions, then  
10 assigning the given entity a work partition selected from said first set of work  
11 partitions.

1 51. (Previously presented) The method of Claim 49 wherein the hierarchy includes odd levels  
2 and even levels, and the method further comprises the steps of assigning work partitions  
3 from odd levels to said first plurality of entities and work partitions from even levels to  
4 said second plurality of entities.

1 52. (Previously presented) The method of Claim 49 wherein performing work partitions in  
2 said first set of work partitions causes said first set of entities produce output consumed  
3 by said second plurality of entities, and performing work partitions in said third set of  
4 work partitions causes said first set of entities to consume output produced by said second  
5 plurality of entities.

1 53-62. (Cancelled)

1 63. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:

4 dividing the operation into a set of work partitions;  
5 assigning work partitions from said set of work partitions to a plurality of entities,  
6 wherein at least one entity of said plurality of entities is assigned a plurality of  
7 work partitions from said set of work partitions; wherein the step of assigning  
8 work partitions is performed by assigning the work partitions in a sequence based  
9 at least in part on sizes associated with the work partitions; and  
10 said plurality of entities operating in parallel on work partitions assigned to them to  
11 perform said operation.

1 64. (Previously presented) The computer-readable medium of Claim 63 wherein the step of  
2 assigning the work partitions in a sequence is performed by assigning relatively larger  
3 work partitions before assigning relatively smaller work partitions.

1 65. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:  
4 dividing the operation into a set of work partitions;  
5 assigning work partitions from said set of work partitions to a plurality of entities,  
6 wherein at least one entity of said plurality of entities is assigned a plurality of  
7 work partitions from said set of work partitions, wherein the step of assigning  
8 work partitions includes  
9 assigning said at least one entity a first work partition from said set of work partitions;  
10 and

11 after said at least one entity has completed operating on said first work partition,  
12 assigning said at least one entity a second work partition from said set of work  
13 partitions;  
14 said plurality of entities operating in parallel on work partitions assigned to them to  
15 perform said operation;  
16 wherein the operation is specified in a query that corresponds to a hierarchy of  
17 operations; and  
18 the step of assigning said at least one entity a second work partition includes  
19 determining whether there are any unassigned work partitions from a first level in  
20 the hierarchy to which said first work partition belonged; and  
21 if there are no unassigned work partitions from the first level in the hierarchy,  
22 then selecting said second work partition from a level in said hierarchy  
23 that is two levels above said first level in said hierarchy.

1 66. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:  
4 dividing the operation into a set of work partitions;  
5 assigning work partitions from said set of work partitions to a plurality of entities,  
6 wherein at least one entity of said plurality of entities is assigned a plurality of  
7 work partitions from said set of work partitions;  
8 said plurality of entities operation in parallel on work partitions assigned to them to  
9 perform said operation;

10 wherein the instructions include instructions for performing the step of generating a serial  
11 execution plan for operations in a database management system (DBMS) running  
12 on a computer system;

13 wherein the instructions include instructions for performing the step of generating a  
14 parallelized execution plan for said serial execution plan, said parallelized  
15 execution plan including first and second operations;

16 wherein the step of dividing an operation is performed by dividing said second operation;

17 wherein the plurality of entities includes one or more slave processes operating on a  
18 plurality of data partitions, the quantity of said data partitions being greater than  
19 the quantity of said slave processes;

20 wherein the instructions include instructions for performing the step of executing said  
21 parallelized execution plan when a plurality of parallel resources of said computer  
22 system are available; and

23 wherein the instructions include instructions for performing the step of executing said  
24 serial execution plan when said plurality of resources are not available.

1 67. (Previously presented) The computer-readable medium of claim 66 wherein said step of  
2 generating a parallelized execution plan includes the steps of:  
3 identifying one or more segments of said serial execution plan that can be parallelized;  
4 and  
5 identifying partitioning requirements of said one or more segments.

1 68. (Previously presented) The computer-readable medium of claim 66 wherein said step of  
2 generating a parallelized execution plan is based on a specification of parallelism in a  
3 statement specifying one of said operations.

1 69. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:  
4 dividing the operation into a set of work partitions;  
5 assigning work partitions from said set of work partitions to a plurality of entities,  
6 wherein at least one entity of said plurality of entities is assigned a plurality of  
7 work partitions from said set of work partitions;  
8 said plurality of entities operating in parallel on work partitions assigned to them to  
9 perform some operation;  
10 generating an execution plan for said operation;  
11 examining said execution plan from bottom up;  
12 identifying a parallelized portion of said execution plan, said parallelized portion can be  
13 processed in parallel, said parallelized portion including first and second  
14 operations, said first and second operations being executable in parallel;  
15 wherein the step of dividing the operation is performed by dividing said second  
16 operation;  
17 wherein the plurality of entities includes one or more slave processes operating on a  
18 plurality of data partitions, the quantity of said data partitions being greater than  
19 the quantity of said slave processes;



20 identifying some serial portion of said execution plan, said serial portion can be  
21 processed in serial; and  
22 allocating a central scheduler between said parallelized portion and said serial portion.

1 70. (Previously presented) The computer-readable medium of Claim 69 further including  
2 instructions for performing the steps of:  
3 identifying a first data flow requirement for a first portion of said execution plan said first  
4 data flow requirement corresponding to a partitioning of a data flow required by  
5 said first portion;  
6 identifying a second data flow requirement for a second portion of said execution plan  
7 said second data flow requirement corresponding by said second portion; and  
8 allocating a data flow director between said first portion and said second portion when  
9 said first data flow requirement is not compatible with said second data flow  
10 requirement said data flow director repartitioning a data flow of said first portion  
11 to be compatible with said second data flow requirement.

1 71. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:  
4 dividing the operation into a set of work partitions;  
5 assigning work partitions from said set of work partitions to a plurality of entities,  
6 wherein at least one entity of said plurality of entities is assigned a plurality of  
7 work partitions from said set of work partitions;

8 said plurality of entities operating in parallel on work partitions assigned to them to  
9 perform said operation;  
10 generating an execution plan to execute database management system (DBMS)  
11 operations in parallel, said execution plan including first and second operations;  
12 wherein the step of dividing said operation is performed by dividing said second  
13 operation;  
14 initiating an operation coordinator in a computer system to coordinate execution of said  
15 execution plan;  
16 initiating, by said operation coordinator, a first set of slaves operating on a plurality of  
17 data partitions to produce data, the quantity of said data partitions being greater  
18 than the quantity of said first set of slave processes;  
19 initiating, as said plurality of entities, by said operation coordinator, a second set of  
20 slaves to consume data; and  
21 directing said second set of slaves to produce data and said first set of slaves to consume  
22 data when said first set of slaves finishes producing data.

1 72. (Previously presented) The computer-readable medium of claim 71 wherein said  
2 execution plan is comprised of operator nodes and said operator nodes are linked together  
3 to form execution sets.

1 73. (Previously presented) A computer-readable medium carrying instructions for  
2 parallelizing an operation, the instructions including instructions for performing the steps  
3 of:

dividing the operation into a set of work partitions;  
assigning work partitions from said set of work partitions to a plurality of entities,  
wherein at least one entity of said plurality of entities is assigned a plurality of  
work partitions from said set of work partitions;  
said plurality of entities operating in parallel on work partitions assigned to them to  
perform said operation;  
generating an execution plan to execute said operations in parallel, said execution plan  
including first and second operations;  
wherein the step of dividing said operation includes dividing said first operation;  
initiating producer slaves operating on a plurality of data partitions to produce a first data  
production;  
initiating consumer slaves to consume said first data production;  
when said first data production is completed, generating an identification of a plurality of  
said consumer slaves that did not receive data in said first data production;  
examining said identification during a subsequent data production; and  
reducing said subsequent data production such that said subsequent data production does  
not produce data for said plurality of said consumer slaves.

74. (Currently amended) A computer-readable medium storing instructions for processing a  
query, the instructions including instructions for performing the steps of:  
receiving a statement that specifies at least an operation;  
determining a user-specified degree of parallelism to use in performing the operation;  
dividing the operation into a set of work partitions;

6 performing a determination of how many entities to use to perform said operation based,  
7 at least in part, on the user-specified degree of parallelism;  
8 assigning work partitions from said set of work partitions to a plurality of entities based  
9 on said determination; and  
10 said plurality of entities operating in parallel on work partitions assigned to them to  
11 perform said operation.

1 75. (Currently amended) The computer-readable medium of Claim 74 wherein:  
2 the query requires a plurality of operations;  
3 the user-specified degree of parallelism is specified in said statement, and  
4 the statement specifies said degree of parallelism for a subset of the plurality of  
5 operations required by the query.

1 76. (Currently amended) The computer-readable medium of Claim 75 wherein  
2 the user-specified degree of parallelism is specified in said statement; and  
3 the degree of parallelism specified by the query indicates that no amount of parallelism is  
4 to be used during execution of a particular portion of the query.

1 77. (Currently amended) The computer-readable medium of Claim 74 wherein  
2 the user-specified degree of parallelism is specified in said statement, and  
3 the degree of parallelism specified by the query indicates a maximum amount of  
4 parallelism to use during execution of said operation.

1 78. (Previously presented) A computer-readable medium carrying instructions for processing  
2 a query, the instructions including instructions for performing the steps of:  
3 dividing an operation required by said query into a set of work partitions by generating a  
4 set of query fragments;  
5 incorporating hints into at least some of said query fragments, wherein the hint associated  
6 with a given query fragment indicates how to perform the work partition  
7 associated with said given query fragment;  
8 assigning query fragments from said set of query fragments to a plurality of entities; and  
9 said plurality of entities operating in parallel on query fragments assigned to them to  
10 perform said operation, wherein entities working on a query fragment associated  
11 with a hint perform the work partition associated with said query fragment in a  
12 manner dictated by said hint.

1 79. (Previously presented) The computer-readable medium of Claim 78 wherein the step of  
2 incorporating hints includes incorporating hints that dictate the operation of a table scan.

1 80. (Previously presented) The computer-readable medium of Claim 79 wherein the step of  
2 incorporating hints that dictate the operation of a table scan includes incorporating hints  
3 that rowid partitioning is to be used during the table scan.

1 81. (Previously presented) The computer-readable medium of Claim 78 wherein the step of  
2 incorporating hints includes incorporating hints that specify performance of a full table  
3 scan.

1 82. (Previously presented) The computer-readable medium of Claim 78 wherein the step of  
2 incorporating hints includes incorporating hints that specify using a particular type of  
3 join.

1 83. (Previously presented) The computer-readable medium of Claim 82 wherein the step of  
2 incorporating hints that specify using a particular type of join includes incorporating hints  
3 that specify using a sort/merge join.

1 84. (Previously presented) The computer-readable medium of Claim 82 wherein the step of  
2 incorporating hints that specify using a particular type of join includes incorporating hints  
3 that specify using a nested loop join.

1 85. (Previously presented) A computer-readable medium carrying instructions for processing  
2 a query, the instructions including instructions for performing the steps of:  
3 determining a hierarchy of operations associated with a query;  
4 dividing a first operation required by said query into a first set of work partitions;  
5 dividing a second operation required by said query into a second set of work partitions,  
6 wherein said second operation immediately follows said first operation in said  
7 hierarchy;  
8 dividing a third operation required by said query into a third set of work partitions,  
9 wherein said third operation immediately follows said second operation in said  
10 hierarchy;

11 assigning work partitions from said first set of work partitions to a first plurality of  
12 entities;  
13 said first plurality of entities operating in parallel on work partitions assigned to them  
14 from said first set of work partitions to perform said first operation;  
15 assigning work partitions from said second set of work partitions to a second plurality of  
16 entities, wherein said second plurality of entities are different entities than said  
17 first plurality of entities; and  
18 said second plurality of entities operating in parallel on work partitions assigned to them  
19 from said second set of work partitions to perform said second operation;  
20 assigning work partitions from said third set of work partitions to said first plurality of  
21 entities; and  
22 said first plurality of entities operating in parallel on work partitions assigned to them  
23 from said third set of work partitions to perform said third operation.

1 86. (Previously presented) The computer-readable medium of Claim 85 further comprising  
2 instructions for performing the following steps when a given entity in said first set of  
3 entities finishes performing a work partition from said first set of work partitions:  
4 determining whether there are any unassigned work partitions from said first set of work  
5 partitions; and  
6 if there are no unassigned work partitions from said first set of work partitions, then  
7 assigning the given entity a work partition selected from said third set of work  
8 partitions; and

9 if there are unassigned work partitions from said first set of work partitions, then  
10 assigning the given entity a work partition selected from said first set of work  
11 partitions.

1 87. (Previously presented) The computer-readable medium of Claim 85 wherein the  
2 hierarchy includes odd levels and even levels, and the instructions further include  
3 instructions for performing the steps of assigning work partitions from odd levels to said  
4 first plurality of entities and work partitions from even levels to said second plurality of  
5 entities.

1 88. (Previously presented) The computer-readable medium of Claim 85 wherein performing  
2 work partitions in said first set of work partitions causes said first set of entities produce  
3 output consumed by said second plurality of entities, and performing work partitions in  
4 said third set of work partitions causes said first set of entities to consume output  
5 produced by said second plurality of entities.

1 89-91. (Canceled).

1 92. (New) The method of Claim 38, wherein the user-specified degree of parallelism is  
2 specified in said statement.

1 93. (New) The method of Claim 38, wherein the user-specified degree of parallelism is  
2 specified for operations that involve a particular table.



1 94. (New) The computer-readable medium of Claim 74, wherein the user-specified degree of  
2 parallelism is specified in said statement.

1 95. (New) The computer-readable medium of Claim 74, wherein the user-specified degree of  
2 parallelism is specified for operations that involve a particular table.